

Preschoolers adapt their exploratory strategies to the information structure of the task

Nora Swaboda (swaboda@mpib-berlin.mpg.de)

MPRG iSearch, Max Planck Institute for Human Development

Azzurra Ruggeri (ruggeri@mpib-berlin.mpg.de)

MPRG iSearch, Max Planck Institute for Human Development
School of Education, Technical University Munich

Zi Lin Sim (zi@berkeley.edu)

Department of Psychology, University of California, Berkeley

Alison Gopnik (gopnik@berkeley.edu)

Department of Psychology, University of California, Berkeley

Abstract

Previous research has shown that active engagement drives children's remarkable learning capabilities. We investigated whether preschoolers are ecological learners, able to select those active learning strategies that are most informative in a given task. Children ($n = 114$; 3 to 5 years old) chose between two exploratory actions (opening vs. shaking) to find an egg shaker hidden in one of four small boxes, contained in two larger boxes. Prior to this game, children learnt that the egg was equally likely to be in any of the four small boxes (Uniform condition), or that it was most likely to be in one particular small box (Skewed condition). Results show that 3- and 4-year-olds, but not 5-year-olds, successfully tailored their exploratory actions to these different likelihood distributions. We suggest that ecological learning may be a key mechanism explaining how children can efficiently learn about the world around them.

Keywords: active learning; ecological learning, adaptiveness, cognitive development

Introduction

How can children learn so much about the world so quickly? A rich body of research has demonstrated that active engagement with the world is a crucial component of learning: Preschoolers (3 to 5 years) and even infants selectively explore objects and events that are likely to be informative, such as those that violate their expectations and assumptions or that are causally confounded (Bonawitz, van Schijndel, Friel, & Schulz, 2012; Cook, Goodman, & Schulz, 2011; Legare, 2012; Schulz & Bonawitz, 2007; Sim & Xu, 2014; Sim & Xu, 2017; Stahl & Feigenson, 2015). As language develops, young children ask about the meaning of words, request the labels of objects, and inquire about the many new and puzzling phenomena they encounter (Chouinard, 2007; Frazier & Gelman, 2009; Mills, Legare, & Grant, 2011; Legare, Mills, Souza, Plummer, & Yasskin, 2013; Ruggeri & Lombrozo, 2015). To learn efficiently, children need to be *ecological learners* (Ruggeri & Lombrozo, 2015; Ruggeri, Sim, & Xu, 2017); that is, they have to *adapt* their active learning strategies to select those

actions or questions that are most informative in the task at hand.

Earlier studies suggest that children do not select the most informative evidence to explore until well into late primary school age (Chen & Klahr, 1999; Kuhn & Brannock, 1977). However, recent studies employing the 20-questions game, in which the goal is to find the solution to a causal inference task (e.g., "Yesterday, Toma was late for school. Why?") by asking only yes/no questions, suggest that children are indeed ecological learners. In particular, 7- to 10-year-olds ask different kinds of questions depending on the likelihood distribution across the potential candidate solutions. For example, children tend to ask questions targeting a single hypothesis (*hypothesis-scanning questions*; e.g., "Is Toma late because he woke up late?") when this hypothesis is presented as more likely than the others. However, when all candidate hypotheses are presented as equally likely, children tend to ask questions that target a feature shared by multiple hypotheses (*constraint-seeking questions*; e.g., "Is Toma late because he could not find something?"), thereby ruling out several candidate hypotheses at once (Ruggeri & Lombrozo, 2015). There is evidence that even 5-year-olds are able to *select* the most informative of two given questions adaptively based on the likelihood distribution in the hypothesis space under consideration - although they are not able to *generate* the most informative questions from scratch yet (Ruggeri, Sim, & Xu, 2017). However, virtually nothing is known about whether children younger than 5 years of age are ecological learners, able to adapt their active learning strategies to the information structure of a task. Moreover, it is unclear whether the adaptiveness found in older preschoolers and primary school children would extend beyond the question-asking paradigm.

The present study

As the question-asking paradigm heavily relies on children's developing verbal skills, it may not be suitable to study ecological learning in younger preschoolers. In the present study, we developed a novel nonverbal exploration paradigm

to investigate whether children as young as three years old are ecological learners. Preschoolers (3 to 5 years old) played a game in which they were presented with two large open boxes, each containing two smaller boxes (Figure 1). They watched the experimenter place an egg shaker in one of the four small boxes and were prompted to retrieve it and use it to activate a light-up toy (*frequency training phase*). This placement was repeated four times. Crucially, we varied between subjects whether the experimenter always placed the egg in the same (Skewed condition) or a different (Uniform condition) small box (Figure 1). Next, children were demonstrated and learned two actions that could be used to find out whether a large box contained the egg: *shaking* it, to hear whether the egg was inside one of the small boxes, or *opening* the large and the two small boxes inside it (*action training phase*).




Figure 1. Frequency training phase. An egg shaker was placed four times into one of four small boxes (green, blue, yellow, and red) contained in two larger boxes (white or black). After each placement, children were asked to retrieve the egg and use it to activate a light-up toy. We manipulated between subjects whether the egg was always hidden in the same (Skewed condition) or a different (Uniform condition) small box.

In the *test phase*, the experimenter asked children to cover their eyes, hid the egg in one of the four small boxes, and closed all the small and both large boxes. The children had to find the egg, so that they could activate the light-up toy. They were allowed to open *only one large box*, but could shake one or both large boxes before deciding which one to open. The shaking and opening actions were informative in different ways in the two conditions. From the frequency training phase, children in the Uniform condition could not know in which small box the egg was hidden. Shaking at least one large box before deciding which one to open was therefore necessary to avoid opening the wrong box and missing the opportunity to find the egg and play with the light-up toy. On the other hand, children in the Skewed condition should have known where the egg was hidden and could open the correct large and small box right away. Drawing a parallel with the question-asking paradigm previously adopted to study ecological learning (Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015; Ruggeri & Lombrozo, 2015; Ruggeri, Sim &

Xu, 2017), shaking a large box in our task corresponds to asking a constraint-seeking question that targets multiple hypotheses (i.e., the two small boxes inside the large box). In contrast, opening a small box corresponds to asking a hypothesis-scanning question that targets one single hypothesis (i.e., one small box inside one large box). If children learned during training how likely the egg was to be found across the four small boxes and adapted their exploratory actions accordingly, they would be more likely to shake before opening a box in the Uniform condition than in the Skewed condition.

Method

Participants. One hundred fourteen children (62 girls, 52 boys) participated in this study (3-year-olds: $N = 38$, $M_{age} = 41.96$ months, $SD = 3.03$ months; 4-year-olds: $N = 37$, $M_{age} = 54.32$ months, $SD = 3.21$; 5-year-olds: $N = 39$, $M_{age} = 64.41$ months, $SD = 3.43$ months). Children were recruited from preschools and museums in the East Bay of the San Francisco area.

Materials. The two large test boxes were made of black and white cardboard. The large boxes could be opened and closed with Velcro strips. Each large test box contained two smaller differently colored round boxes (green/blue for the white large box and yellow/red for the black large box), glued to its bottom. To ensure that the small boxes were easy to open, plastic rings were attached to their lids. The two large boxes used for the action training were identical in size and shape to the two large test boxes but were made of brown and grey cardboard. The small boxes inside the large box used to demonstrate the opening action were purple and pink. Two white egg shakers were used: one with red dots for the frequency training phase and the test phase, and one with blue dots for the action training phase. The spinning light-up toy was mounted on top of a black box and was surreptitiously activated via remote control by the experimenter whenever the egg shaker was placed on the machine. We extensively piloted the materials and procedure with children aged 2 to 6 years to make sure that the size and weight of the large boxes, relevant for the shaking action, as well as their opening mechanism, were appropriate for all children, so that even the youngest ones could perform the shaking and opening actions without difficulty.

Design and Procedure. We employed a between-subjects design with two conditions (Skewed vs. Uniform) and three age groups (3, 4, and 5 years). The experiment consisted of four phases. During the *familiarization phase*, the experimenter presented the egg shaker and placed it on the light-up toy to demonstrate that the egg could be used to activate the toy. She then told the children, “In this game, when you find this ball, you’ll get to put it on the toy and play with it!” Next, she showed children the two large boxes, which sat on the table with their lids open. She then opened and closed the four small boxes one after the other to demonstrate that they were empty.

The *frequency training phase* consisted of four rounds. On each round, the experimenter opened one of the four small boxes, placed the egg shaker inside it, and closed it again. In the Skewed condition ($N = 56$), the experimenter always placed the egg shaker in the same small box, either the leftmost or the rightmost (counterbalanced across participants), saying, “See? In this game, I always hide the ball in the same box” (see Figure 1). In the Uniform condition ($N = 58$), the experimenter placed the egg shaker in a different small box on each round in an ordered fashion, either right-to-left or left-to-right (counterbalanced across participants), so that at the end of the training phase the experimenter had placed the egg shaker once in each small box, saying, “See? In this game, I always hide the ball in a different box” (see Figure 1). After the egg shaker had been placed inside one of the four small boxes, the children were prompted to retrieve the egg shaker and use it to activate the light-up toy.

During the *action training phase*, the experimenter demonstrated two actions the children could perform to find out whether or not a large box contained the egg shaker: they could either *shake* the large box, to hear if one of the small boxes inside contained the egg shaker, or *open* the large and the small boxes to explore their contents. Each action was demonstrated using novel large and small boxes (see Materials section above). For the shaking training, the experimenter presented a novel large box and said, “Oh look, here is another big box. This big box also has two small boxes inside, like those [points to the two test boxes]. Hmm... I wonder if there is a ball in one of the two small boxes inside here.” After a short moment, she continued, “Here is one thing we can do to find out: we can shake the big box!” She then proceeded to shake the box, ensuring that the children had noticed the egg shaker inside it from its rattling sound. Children were then prompted to shake the box themselves.

For the opening training, the experimenter presented the other novel large box and said, “Oh look, here is another big box. This big box also has two small boxes inside, like those [points to the two test boxes]. Hmm... I wonder if there is a ball in one of the two small boxes inside here.” She continued, “Here’s one thing we can do to find out: we can open the big box!” The experimenter then demonstrated how to open the large box and the two small boxes inside it and prompted children to repeat the action. The order in which the two actions were demonstrated was counterbalanced across participants.

During the *test phase*, the experimenter presented children with the same boxes used in the frequency training phase and said, “Okay, now we are going to play a hiding game! I’m going to hide the ball in one of the small boxes, and you’ll have to find it!” Children were then asked to cover their eyes with the help of their parents. To hide the egg shaker, the experimenter opened the lids of all four small boxes, placed the egg shaker in the target box, and closed all four small boxes again but left the lid of the large boxes open. Opening and closing of the four small lids was performed in random order. The egg’s hiding location varied so that in the Uniform condition, the egg shaker was hidden in a random small box.

In the Skewed condition, it was placed either in the leftmost or rightmost small box, depending on the current counterbalancing version.

After hiding the egg shaker, the experimenter pushed the large boxes to the edges of the table so that they would be hard to reach for the children and said, “I’m ready now! You can look! I hid the ball in one of the four small boxes, that are now over there [indicates the boxes on the table]. When you find it, you can play with the toy!” Next, she closed both large boxes and told the children that they were allowed to open *only one* of the two large boxes to find the egg shaker.

Results

To analyze children’s performance in our experiment we performed a series of Chi-square tests. Overall, we found that more children (67%) shook one or both large boxes before deciding which one to open in the Uniform as compared to the Skewed condition (43%), $\chi^2(1, 114) = 6.85, p = .009$. This suggests that the children had effectively learned how likely the egg was to be found across the four small boxes and had used this information to decide which exploratory action to perform.




Figure 2. Proportion of children who shook one or both large boxes before deciding which one to open in the Uniform and Skewed conditions. Three- and 4-year-olds efficiently adapted their exploratory actions to the task: they were more likely to shake a large box before opening it in the Uniform than in the Skewed condition. Five-year-olds were equally likely to shake first, irrespective of condition.

We observed a general developmental difference in the actions children performed: across both conditions, most of the 3-year-olds opened one large box right away (66%), whereas most of the 4- (59%) and 5-year-olds (72%) shook first ($\chi^2(2, 114) = 11.39, p = .003$; see Figure 2). Our results also showed a decrease in adaptiveness with age: As can be seen in Figure 2, 3- and 4-year-olds performed different exploratory actions (shaking vs. opening) in the two conditions, with a larger proportion of children shaking

before opening in the Uniform as compared to the Skewed condition (3-year-olds: $\chi^2(1, 38) = 4.68, p = .03$; 4-year-olds: $\chi^2(1, 37) = 4.89, p = .03$). However, this was not the case for 5-year-olds ($p = .65$), the majority of whom shook first in both conditions. This finding seems to suggest that whereas 3- and 4-year-olds are ecological, adaptive learners, 5-year-olds might not be sensitive to the task's information structure.

Discussion

In this project, we investigated whether preschoolers adapt their active learning and search strategies to the information structure of the current learning environment. Children searched for an egg shaker hidden in one of four small boxes contained in two larger boxes. In two different conditions, children initially learned that the egg shaker could be found in any of the four small boxes (Uniform condition) or was most likely to be found in one particular small box (Skewed condition). Children could then choose between two exploratory actions: they could either open the large and small boxes right away, or shake them first to hear which large box contained the egg. We found that children in the Uniform condition tended to shake the boxes first, whereas children in the Skewed condition preferred to open a box right away. In line with previous research demonstrating that preschoolers and even infants are excellent at tracking statistical regularities (e.g., Denison, Bonawitz, & Gopnik, 2013; Kushnir, Xu, & Wellman, 2010; Waismeyer, Meltzoff, & Gopnik, 2015; Wellman, Kushnir, Xu, & Brink, 2016), our results suggest that preschoolers in our task correctly inferred how likely the egg was to be found across the four small boxes based on the frequency pattern they observed.

More importantly, we demonstrated that preschoolers exploit this statistical sensitivity to guide their own exploratory actions – choosing the action that promises the largest information gain in the current learning situation. We have thus provided compelling evidence that even 3-year-olds are ecological learners, able to efficiently adapt their exploratory actions to the information structure of the task. This result is particularly important considering that earlier investigations have offered mixed evidence as to whether younger preschoolers can implement the most efficient learning strategies when asking questions from scratch or selecting among given questions (Herwig, 1982; Legare et al., 2013; Ruggeri, Sim, & Xu, 2017).

We also found a developmental increase in the percentage of children who shook one or both large boxes, suggesting that with increasing age, children's default strategy changed from opening to shaking. Interestingly, this shift is consistent with the developmental progression from hypothesis-scanning to constraint-seeking questions that occurs later in development (Mosher & Hornsby, 1966; Ruggeri & Lombrozo, 2015; Ruggeri, Sim, & Xu, 2017). To test the hypothesis that this pattern reflects a change in "default strategy", we are currently analyzing children's response latencies, defined here as the time elapsed between the experimenter's "go" signal and the child's first touch of a box. If children's default strategy changes from opening to

shaking, we might expect younger children to be faster to select the opening than shaking action, and older children to be faster to select the shaking action.

In addition, we are also assessing age related differences in the number of children who opened the correct *small* box in the two conditions. Generally, we expect fewer children in the Skewed condition to open the wrong small box, as compared to the Uniform condition, where we expect chance performance. One might argue that those children who opened the wrong small box in the Skewed condition did not actually perform the appropriate action. However, it could also be that children who opened the wrong small box did encode the skewed nature of the likelihood distribution correctly but failed to accurately remember which of the small boxes was the one containing the egg shaker. Considering that working memory improves significantly over the preschool years (e.g. Roman et al., 2014), we might expect a developmental decrease in the number of children opening the wrong small box in the Skewed condition.

We also observed a surprising developmental change in adaptiveness: Whereas 3- and 4-year-olds tended to select the exploratory action that was most efficient in the two conditions (i.e., shaking in the Uniform and opening in the Skewed condition), the choices of 5-year-olds were not adaptive. Although counter-intuitive, this finding is in fact in line with previous research showing that despite a general increase in search efficiency across the lifespan (Ruggeri & Lombrozo, 2015; Ruggeri & Feufel, 2015), the *adaptiveness* of search strategies may decrease in adulthood. For instance, 9-year-old children, but not adults, asked different types of questions depending on the likelihood of the solution in a 20-questions game where they had to find out why John was late for work. In particular, when told that the solution to the game was unlikely (i.e., infrequent), children asked more constraint-seeking questions than when they were told the solution was very likely. However, adults always asked a majority of constraint-seeking questions, irrespective of the solution's likelihood (see Ruggeri & Lombrozo, in preparation). Similarly, evidence from causal learning studies suggests that although overall younger children are less efficient learners, they might be more *sensitive* than adults to the evidence they observe, especially when learning about unusual causal systems (Lucas, Bridgers, Griffiths, & Gopnik, 2014).

How can the decrease in adaptiveness in our task be explained? There are several not mutually exclusive possible interpretations. First, it could be that growing up, children learn that the world is more likely to be ruled by uncertainty than by clearly predictable risk patterns. Hence, 5-year-olds may have mistrusted the deterministic model of the world we presented them in the Skewed condition, in which the egg shaker was placed in the same small box every time. Second, a more enhanced understanding of the social world, including others' beliefs, intentions and the possibility of deception (e.g., Bosco & Gabbatore, 2017; Lee, 2013; Polak & Harris, 1999; Wellman & Cross, 2001) may have led older preschoolers to distrust the experimenter to repeat the pattern

presented during training. Instead, children may have suspected her to be tricking them by inducing a false belief about the egg's hiding location. Crucially, although from a purely information theoretical perspective shaking the large boxes in the Skewed condition was an unnecessary action, it was not penalized in any way. As a result of suspicion towards our experimental set-up, older children may have decided to shake the large boxes because they preferred to sacrifice efficiency in favor of certainty – preferring to ensure, at no cost, that the egg was where expected.

To test the hypothesis that older children may have been more suspicious or mistrusting of our experimental set-up, we are conducting further analyses on our data, assessing developmental differences in the proportion of children who decided to shake one or both boxes before opening. Indeed, in both conditions, whatever intuitions children might have had about the location of the hidden egg, shaking the second large box did not provide any additional information. Such a behavior would therefore be a sign of confirmatory tendencies, reflecting children's mistrust in the experimenter ("Did she *really* hide the egg at all?"). Thus, we might expect older children to be more likely to shake more than one large box.

Finally, older preschoolers may have experienced more often than younger children that constraint-seeking strategies are generally very effective in situations of uncertainty and may have implemented such a strategy (i.e., shaking) as a default without considering the characteristics of the task. To rule out some of these factors, we are currently conducting a follow-up study in which we motivate 5-year-olds to shake a large box only when strictly necessary by asking them to "pay" a sticker to do so. In addition, we plan to present children with a more realistic model of the world by introducing several probabilistic versions of the Skewed condition. In this version of the task, we will gradually vary the likelihood of the egg shaker being found in the different small boxes, providing more fine-grained differences in informativeness between the shaking and opening actions. Finally, we plan to replicate our findings in different cultural and educational contexts to establish the robustness and universality of active learning adaptiveness. In fact, there is a large number of sociocultural factors that may impact children's behavior in an active learning task like the one used in this study. For instance, growing up in a more or less performance-oriented environment may influence whether children interpret our game as a performance-test, in which opening the wrong box indicates *failure*, or whether they think of it as just a game where performance has no implications. Similarly, it may be interesting to explore how schooling affects children's ecological learning potential. For example, do educational programs focusing on nurturing children's natural exploration tendencies, such as Montessori schools, boost their active learning effectiveness and adaptiveness? Investigating the impact of these factors might help to develop interventions and to design intuitive environments to improve children's active and ecological learning at different developmental stages.

The present paper provides, for the first time, compelling evidence that children as young as three are ecological learners who adapt their exploratory actions to a task's information structure. The efficiency of a learning strategy cannot be measured in absolute terms because strategies vary in informativeness depending on the characteristics of the task at hand (Todd, Gigerenzer, & The ABC-Research Group, 2012). In this sense, the ability to adapt learning strategies to characteristics of the environment is crucial to ensure learning effectiveness. Indeed, ecological learning provides a key mechanism underlying children's remarkable learning capacities: We suggest that children are excellent learners *because* they are able to flexibly tailor and adapt their exploratory actions to characteristics of the current task and dynamically choose those actions that maximize information gain. Therefore, it is crucial to investigate more closely how the potential for ecological learning develops across the lifespan and which factors have an influence on it.

Acknowledgments

We thank Patricia Finnegan, Kritika Shrestha, Sarah De La Vega, Gwendolin Ngoh, and Anqi Li for assistance in recruiting, data collection, and coding. We would also like to thank the Bay Area Discovery Museum in San Francisco as well as the participating preschools in Berkeley for providing research space.

References

- Bonawitz, E. B., van Schijndel, T. J., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology, 64*, 215-234.
- Bosco, F. M., & Gabbatore, I. (2017). Sincere, deceitful, and ironic communicative acts and the role of the Theory of Mind in childhood. *Frontiers in Psychology, 8*.
- Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child Development, 70*, 1098-1120.
- Chouinard, M. M. (2007). Children's questions: A mechanism for cognitive development. *Monographs of the Society for Research in Child Development, 72*, 1-108.
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition, 120*, 341-349.
- Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition, 126*, 285-300.
- Frazier, B. N., Gelman, S. A., & Wellman, H. M. (2009). Preschoolers' search for explanatory information within adult-child conversation. *Child Development, 80*, 1592-1611.
- Herwig, J. E. (1982). Effects of age, stimuli, and category recognition factors in children's inquiry behavior. *Journal of Experimental Child Psychology, 33*, 196-206.
- Kuhn, D., & Brannock, J. (1977). Development of the isolation of variables scheme in experimental and "natural

- experiment" contexts. *Developmental Psychology*, *13*, 9-14.
- Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, *21*, 1134-1140.
- Lee, K. (2013). Little liars: Development of verbal deception in children. *Child Development Perspectives*, *7*, 91-96.
- Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development*, *83*, 173-185.
- Legare, C. H., Mills, C. M., Souza, A. L., Plummer, L. E., & Yasskin, R. (2013). The use of questions as problem-solving strategies during early childhood. *Journal of Experimental Child Psychology*, *114*, 63-76.
- Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*, *131*, 284-299.
- Mills, C. M., Legare, C. H., Grant, M. G., & Landrum, A. R. (2011). Determining who to question, what to ask, and how much information to ask for: The development of inquiry in young children. *Journal of Experimental Child Psychology*, *110*(4), 539-560.
- Mosher, F. & Hornsby, J.R. (1966). On asking questions. *Studies in Cognitive Growth*, 86-102.
- Polak, A., & Harris, P. L. (1999). Deception by young children following noncompliance. *Developmental Psychology*, *35*(2), 561.
- Roman, A. S., Pisoni, D. B., & Kronenberger, W. G. (2014). Assessment of working memory capacity in preschool children using the missing scan task. *Infant and Child Development*, *23*(6), 575-587.
- Ruggeri, A., & Feufel, M. A. (2015). How basic-level objects facilitate question-asking in a categorization task. *Frontiers in Psychology*, *6*, 918.
- Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, *143*, 203-216.
- Ruggeri, A., & Lombrozo, T. (2018) Children, but not adults, adapt their questions to the likelihood of the solution, *in preparation*.
- Ruggeri, A., Sim, Z. L., & Xu, F. (2017). "Why is Toma late to school again?" Preschoolers identify the most informative questions. *Developmental Psychology*, *53*, 1620.
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental psychology*, *43*, 1045.
- Sim, Z., & Xu, F. (2014). Acquiring Inductive Constraints from Self-Generated Evidence. In P. Bello, M. Guarini, M. McShane, & S. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 1431-1436). Austin, Texas: Cognitive Science Society.
- Sim, Z. L., & Xu, F. (2017). Learning higher-order generalizations through free play: Evidence from 2-and 3-year-old children. *Developmental Psychology*, *53*, 642.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*, 91-94.
- Todd, P. M., Gigerenzer, G. & The ABC-Research Group (2012). *Ecological rationality: Intelligence in the world*. OUP USA.
- Waismeyer, A., Meltzoff, A. N., & Gopnik, A. (2015). Causal learning from probabilistic events in 24-month-olds: an action measure. *Developmental Science*, *18*, 175-182.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: the truth about false belief. *Child Development*, *72*(3), 655-684.
- Wellman, H. M., Kushnir, T., Xu, F., & Brink, K. A. (2016). Infants use statistical sampling to understand the psychological world. *Infancy*, *21*, 668-676.